

## Background

The Torah Academy of Bergen County (TABC) is an educational institution, which is becoming more and more dependent on services that are supported via a network. Email and file servers are good examples of services that are supported over a network. For email to work over a network, there is a central computer running a program that allows authorized users access to their personal email account and permits them to view and send messages. The user account could be accessed from any computer within the school or via the Internet.

File servers have directories setup on a central computer that allows people to save, view, and update files that are of interest to others in the school. Sharing a class calendar is just one example of a file that would be available via the file server service.

It is not surprising to note that files (data) stored on a computer might be lost – meaning that the data cannot be retrieved. There are a number of reasons for this - the motor that is spinning the disk drive might malfunction, a virus might infect the computer, or, a person might delete a file by accident. In all of these situations, it will take time to restore the central computer to a state where it will be accessible to everyone. To restore a deleted file could take a couple of hours while replacing a disk drive or cleaning an infected computer might take a few days. In both of these cases, TABC doesn't have the resources available to address either of these issues. For example, the CBL Data Recovery Company

can copy 20G (gigs) of data from a damaged disk drive and ship the recovered data within twenty-four hours at the cost of approximately one thousand two hundred dollars.

Losing the services of a central server for as little as an hour presents an organization such as TABC with tangible and intangible expenses created by the downtime. To help reduce these costs, organizations typically implement fail-over systems, where if one server goes down another server picks up the former's responsibilities. For example, if the power supply on one email server overheats and shuts down, another server will take over providing email services. This setup tries to prevent hardware failures from causing a lapse in network services.

The fail-over concept has proven to be very reliable, but it is expensive to implement and maintain. It requires a set of computers connected together forming what is known as a cluster. Each computer in the cluster must be a clone of each other – they must have the same hardware and software installed since if one of the computers stops working, any other computer in the cluster could take its place. A cluster also requires specialized hardware to support the many users and the fail-over concept – a large array of hard disks. Having the array of disks available via the network and separate from any one computer within the cluster guarantees that when a fail-over occurs, there is no question that the data stored on those disks is current. For example, since each computer, or node, of the cluster uses the network storage to access email information, if one node goes down and another node starts up – the new node

will provide the same email data from the network storage device as the node that just went down.

Many of the utilities used to set up such a cluster and monitor performance are designed with the expectation an experienced system administrator and/or a network specialist is going to use them. The companies that have clustering solutions on the market today assume that the institutions considering their clustering solution have a full time network staff that is familiar with many of the concepts of network administration. These companies have provided a large number of options that make their clustering solutions very powerful – and the administrator must be knowledgeable enough to know what options to use and how to configure them.

Small business and schools, such as TABC, are not exempt from the responsibility of providing reliable network services to their employees/students. TABC could benefit immensely from the stability a fail-over cluster affords, but they cannot afford the expensive hardware and software required to implement any of the existing commercial fail-over clusters. The problem is to meet the reliability needs of small organizations by implementing an easy to install and simple to use fail-over clustering system while keeping the costs of the project to the bare minimum.